

Probabilistic Selection in AGENTSPEAK(L)

Francisco Coelho
fc@di.uevora.pt

Vitor Nogueira
vbn@di.uevora.pt

Abstract

Agent programming is mostly a symbolic discipline and, as such, draws little benefits from probabilistic areas as machine learning and graphical models. However, the greatest objective of agent research is the achievement of autonomy in dynamical and complex environments — a goal that implies embracing uncertainty and therefore the entailed representations, algorithms and techniques. This paper proposes an innovative and conflict free two layer approach to agent programming that uses already established methods and tools from both symbolic and probabilistic artificial intelligence. Moreover, this framework is illustrated by means of a widely used agent programming example, GOLDMINERS.

1 INTRODUCTION

Agent autonomy is a key objective in Artificial Intelligence (AI). Complex and dynamic environments, like the physical world where robots must delve, impose a degree of uncertainty that challenges the vocation of symbolic processing. But while a probabilistic approach — currently expressed in Machine Learning (ML) and Probabilistic Graphical Models (PGMs) [1]— is required for certain aspects of such tasks, a great deal of agent programming is better handled by declarative programming (*e.g.* PROLOG) and more specifically, Beliefs, Desires and Intentions (BDI) architectures for autonomous agents, part of symbolic AI.

2 STATE OF THE ART

Although the symbolic and probabilistic areas of AI correspond, in fact, to different and often antagonist cultures and perspectives, they are not necessarily incompatible. Bridges between them are being built based on distribution semantics [2] or markov random fields [3]. From that common ground there are two possible paths towards the interplay of symbolic and probabilistic AI: the extension of PGMs with logical and relational representations (done by Statistical Relational Learning (SRL)) [4, 5] and the extension of logic programming languages with probability, in Probabilistic Logic Programming (PLP) [6, 7, 8].

From the point of view of programming autonomous agents the symbolic/probabilistic division essentially still persists: symbolic architectures, such as BDI, describe the behaviour of the agents on the basis of metaphors (*e.g.* goals, beliefs) drawn from human behaviour while the principle of Maximum Expected Utility (MEU) is included, as influence diagrams, in probabilistic AI but there is only seminal work blurring that division.




Figure 1: The JASON deliberation process very resumed, with selection functions highlighted. The ASL only specifies the signature of the functions omitting any conditions, besides the type, on the output.

Concerning agents programming JASON [9] is a popular AGENTSPEAK(L) (ASL) [10] interpreter and framework, triggering a considerable amount of research (*e.g.* [11, 12]). The BDI architecture in general, including ASL and JASON in particular, outline a set of symbolic data structures and processes with more or less detailed semantics. The ASL as implemented in JASON specifies that the deliberation cycle, depicted in figure 1, certain selection steps are handled by certain functions. It also defines the signatures of these functions but omits their inner workings. Such omissions play a central role in this work. Despite some work concerning intention selection [13] the default selection function implementation in JASON is a simple process based in round-robin scheduling: intentions form a stack and at each time-step the head action of the top intention is selected; that intention is then sent to the bottom of the stack. This somewhat simplistic approach to selection is good enough for many tasks, including winning planning competitions [14, 15].

However we can see JASON agents in trouble when their environment becomes stochastic. This assertion is hinted by a simple experiment plotted in figure 2: the GOLDMINERS is a virtual scenario used in the 2006 Multi-Agent Programming Contest [17] edition, now part of JASON’s examples. If this scenario is run unchanged the two playing teams reach scores that are clearly reduced if even a small amount of noise is added (5% in the plotted experiment) to the perceptions and actions.

It turns out that a Bayesian Network (BN) is a natural representation of the complex interdependency of random variables and, therefore, a great candidate to represent probabilistic beliefs. But the task of replacing symbolic beliefs by BNs is far from trivial in part because changing the data structure of beliefs entails a chain of reconsiderations about every aspect of the architecture. For example, plans have contexts that must be unifiable with the agent’s beliefs base; changing the beliefs base from a set of closed formulas to a joint distribution of random variables will break (unchanged) unification with those contexts.

3 PROBABILISTIC SELECTIONS

Currently the problem of extending ASL data structures with probabilistic features is being addressed by different authors [18, 19, 20, 21, 22] but isn’t yet fully solved. An alternative and less intrusive application of probabilistic AI to ASL targets the processes instead of the data. Bounding probabilistic techniques to the computation of ASL selection functions (events, options and intentions¹, as in figure 1) promises a

¹Recent versions of JASON include an inbox, an outbox and a select message function.




Figure 2: If sensors report misreadings of the environment state the symbolic inference process inherent to BDI uses false perceptions as (true) facts of the environment and the deliberation process works on wrong assumptions. This plot relates sensor noise (the rate of sensor misreadings, in the horizontal axis) with agent performance measured by the the number of gathered golds (in the vertical axis). Two teams are plotted, the basic reference “dummy” that barely uses BDI features and the “smart” team, fully BDI, (designed by [15]) that won the 2006 “Multi-agent Programming Contest” [16] featuring the GOLDMINERS scenario. Each data point summarizes the number of gathered golds by team in a given noise parameter and consists of the mean and standard variation of ten samples. The mean is traced by a thin black line and standard variation by a band centred in the mean value. Values between data-points are interpolations.

number of benefits:

- selection functions usually have natural formulations in terms of optimization problems which, in many cases, are well handled by probabilistic algorithms;
- since their computations are unspecified (in ASL or JASON) probabilistic techniques can be used without compromising previous work;
- symbolic and probabilistic AI roles are clearly separated but both simultaneously contribute to the agent behaviour:
 - symbolic programming uses unchanged ASL to define high level agent behaviour, with plans, goals, beliefs, resolution,*etc*;
 - probabilistic algorithms use unchanged tools (*e.g.* WEKA [23, 24, 25, 26] or SAMIAM²) to process low level noisy signals — with BNs, influence diagrams, monte-carlo markov chains, expectation-maximization, *etc*;

In summary, defining the selection functions of ASL as optimization tasks to be solved by probabilistic techniques seems to pose a promising set of open problems, with the potential of great contribute for an evolution (and not a revolution) of agent programming.

²From <http://reasoning.cs.ucla.edu/samiam/>.

3.1 PROBLEM STATEMENT

Our immediate goal is to use the original ASL JASON programs of the miners team, the stochastic environment of the experiment depicted in figure 2 and define the computation of the intention selection with a probabilistic process based in the MEU principle. Success achieving this goal can be directly measured by the effect on the team’s performance: the greater the performance increase, the greater the success. This paper motivates and outlines the design of such probabilistic process. Implementation and evaluation are postponed to future developments.

Next a more precise statement concerning intention selection in JASON is preceded by short overviews of the GOLDMINERS competition, noise model used, miners coordination and influence diagrams.

The original environment of the GOLDMINERS competition is partially observable, stochastic, sequential, dynamic and discrete (see [27] about this classification). The competition description states that “[t]he environment of the multi-agent system was a grid-like world where agents could move from one cell to a neighbouring cell if there was no agent or obstacle already in that cell. In this environment, gold could appear in the cells. Participating agent teams were expected to explore the environment, avoid obstacles and compete with another agent team for the gold. The agents of each team could coordinate their actions in order to collect as much gold as they could and to deliver it to the depot where the gold can be safely stored. Agents had only a local view on their environment, their perceptions could be incomplete, and their actions could fail [*in* [16]].”

Although noise is present in the original competition scenario in the form of incomplete perception and action failure the GOLDMINERS examples in the JASON distribution (besides providing a proxy to that competition simulator) optionally use a local simulator for development and evaluation purposes. In that local simulator noise only increases the probability of action failure in proportion to current cargo. Also the environment type changes from stochastic (where gold can appear in cells) to strategic (because the only changes in the environment state are produced by agents). Experimental results depicted in figure 2 result from this local simulator with added noise in the perceptions (number of gold pieces, *etc.*) and action choice. The amount of noise is configurable by an external parameter that when zero the simulation becomes the original and local, noise free, one. The GOLDMINERS miners team follows a coordination protocol concerning the collection and transport of newly found gold pieces. There are two kinds of agents, a single leader and a set of miners. “[The] leader helps the miners to coordinate themselves in [...] the negotiation process that is started when a miner sees a piece of gold and is not able to collect it (because its container is full). This miner broadcasts the gold location to other miners who then send bids to the leader. The leader chooses the best offer and allocate the corresponding agent to collect that piece of gold [*in* [15, 12]].”

The default JASON intention selection method, round-robin based, is not much context aware. While extensions like plans “priority” annotations can provide cues to more informed choices, action selection is the subject of an huge area of probabilistic AI centred around the MEU principle. Within the PGM setting this principle is instantiated by influence diagrams, graphical models extended with special nodes to represent utilities and actions.




Figure 3: An influence diagram for intuition selection of GOLDMINERS agents. The discovery, auctioning, transportation and deposit of gold pieces of the coordination protocol is represented in the graph. Variables (denoted by round rectangles) and Utilities (chamfered rectangles) in this diagram refer to *total quantities* of gold pieces. The probabilistic effect of the selected action is represented by the arrows that leave the action (rectangular) node.

3.2 RESOLUTION PATH

Given the GOLDMINERS scenario the miners intention selection function is described using an influence diagram as depicted in figure 3. Deposited golds define an utility node and the range of actions is extracted from active intentions. The miners team coordination protocol is represented from the start with the discovery of a new gold piece to termination with the deposit of that gold piece.

The resulting utility function can then be used by the MEU principle to select, from the available action (the heads of instantiated plans in the intentions stack), the optimal one. Nodes in the influence diagram represent Conditional Probability Distributions (CPDs) that, as part of the resolution proposed here, are to be tuned. Once the influence diagram defined, existing java PGM libraries that supports influence diagrams (e.g. SAMIAM) generate a (static) deliberation policy that can be inserted into the agents intention selection function defined in JASON.

4 CONCLUSION

Looking forward the GOLDMINERS exercise two issues require further consideration:

- “real-time” and “off-line” symbolic/probabilistic levels interplay;
- utility/goal conflicts;
- multi-agent applications;

In the deliberation process the communication between symbolic and probabilistic levels is of key importance. Selection functions signatures already define some “real-time” channels: arguments carry information from the symbolic to the probabilistic level while the returned values work the other way around. But perhaps other forms of mutual “off-line” influence can be considered. For example, a process where at the probabilistic level influence diagrams (e.g. BN) are structured from the set of symbolic plans and beliefs and the coefficients in the factors of the nodes are learned from experience. One can also imagine a similar influence going from the probabilistic to the

symbolic level, either introducing relevant concepts from unsupervised feature learning [28] using Deep Neural Networks (DNN) [29, 30] or summarizing ML techniques like Principal Component Analysis (PCA) [31].

The two-layer approach also raises some practical and theoretical concerns on behaviour consistency. Since the probabilistic level maximizes utility functions and the symbolic level derives intentions from goals one as to ask how the utilities relate with the goals: if they are inconsistent the agent will be in trouble like HAL 9000, as explained in the movie “2010”.

Multi-agent systems imply social communication that in JASON is represented by messages that must be selected from an in-box and composed and sent to the out-box. Here the probabilistic layer can contribute to the social perception in other ways besides the computation of the select message function. For example, by using BNs to represent reputation and expectations about other agents then bayesian learning updates the social reasoning.

Coupling symbolic and probabilistic AI is an hard task. Our approach tries to minimize and loosely regulate contact by the separation of competences. Such separation seems easier in a structured framework as JASON than in a (somewhat) simpler but much broader setting as PGM or PLP. The proposed approach seems technically feasible and with the potential to lead to a synergy in both flavours of AI.

Future work

- Apply the two-layer approach outlined here in a virtual scenario (eg GOLDMINERS);
- Consider probabilistic correction on other elements of the BDI process. In particular, as a pre-process of the Belief-Revision function;
- Formulate a probabilistic version BDI using SRL or PLP;

ACKNOWLEDGEMENTS

The people around us, the flow of experiences, the internet.

References

- [1] D. Koller and N. Friedman, *Probabilistic graphical models: principles and techniques*. The MIT Press, 2009.
- [2] T. Sato, “A statistical learning method for logic programs with distribution semantics,” in *Proceedings of the 12th International Conference on Logic Programming (ICLP’95)*, Citeseer, 1995.
- [3] P. Domingos, S. Kok, H. Poon, M. Richardson, and P. Singla, “Unifying logical and statistical AI,” in *AAAI*, vol. 6, pp. 2–7, 2006.
- [4] T. Sato, “A general MCMC method for bayesian inference in logic-based probabilistic modeling,” in *Proceedings of the Twenty-Second international joint conference on Artificial Intelligence-Volume Volume Two*, pp. 1472–1477, AAAI Press, 2011.

- [5] T. Sato, “Inside-outside probability computation for belief propagation.,” in *IJ-CAI*, pp. 2605–2610, 2007.
- [6] D. Fierens, G. V. den Broeck, J. Renkens, D. Shterionov, B. Gutmann, I. Thon, G. Janssens, and L. D. Raedt, “Inference and learning in probabilistic logic programs using weighted boolean formulas,” *arXiv*, p. 1304.6810, 04 2013.
- [7] D. Fierens, G. V. d. Broeck, I. Thon, B. Gutmann, and L. De Raedt, “Inference in probabilistic logic programs using weighted CNF’s,” *arXiv*, p. 1202.3719, 2012.
- [8] B. Gutmann, I. Thon, and L. De Raedt, “Learning the parameters of probabilistic logic programs from interpretations,” in *Machine Learning and Knowledge Discovery in Databases*, pp. 581–596, Springer, 2011.
- [9] R. H. Bordini, J. F. Hübner, and M. Wooldridge, *Programming multi-agent systems in AgentSpeak using Jason*. Wiley, 2007.
- [10] A. S. Rao, “AgentSpeak (L): BDI agents speak out in a logical computable language,” in *Agents Breaking Away*, pp. 42–55, Springer, 1996.
- [11] R. H. Bordini and J. F. Hübner, “Semantics for the Jason variant of AgentSpeak (plan failure and some internal actions).,” in *ECAI*, pp. 635–640, 2010.
- [12] R. H. Bordini and J. F. Hübner, “BDI agent programming in AgentSpeak using Jason,” in *Computational logic in multi-agent systems*, pp. 143–164, Springer, 2006.
- [13] R. H. Bordini, A. L. Bazzan, R. de O Jannone, D. M. Basso, R. M. Vicari, and V. R. Lesser, “AgentSpeak (XL): Efficient intention selection in BDI agents via decision-theoretic task scheduling,” in *Proceedings of the first international joint conference on Autonomous agents and multiagent systems: part 3*, pp. 1294–1302, ACM, 2002.
- [14] R. H. Bordini, J. F. Hübner, and D. M. Tralamazza, “Using Jason to implement a team of gold miners,” in *Computational Logic in Multi-Agent Systems*, pp. 304–313, Springer, 2007.
- [15] J. F. Hübner and R. H. Bordini, “Developing a team of gold miners using jason,” in *Programming Multi-Agent Systems*, pp. 241–245, Springer, 2008.
- [16] M. Dastani, J. Dix, and P. Novák, “The second contest on multi-agent systems based on computational logic,” in *Computational Logic in Multi-Agent Systems*, pp. 266–283, Springer, 2007.
- [17] T. Behrens, J. Dix, M. Köster, and J. Hübner, “Special issue about multi-agent-contest II,” *Annals of Mathematics and Artificial Intelligence*, vol. 61, 2011.
- [18] B. Luz, F. Meneguzzi, and R. Vicari, “Alternatives to threshold-based desire selection in bayesian BDI agents,” in *1st International Workshop on Engineering Multi-Agent Systems*, pp. 208–223, 2013.
- [19] M. S. Fagundes, R. M. Vicari, and H. Coelho, “Deliberation process in a BDI model with bayesian networks,” in *Agent Computing and Multi-Agent Systems*, pp. 207–218, Springer, 2009.

- [20] M. S. Fagundes, “Integrating BDI model and Bayesian Networks,” Master’s thesis, Universidade Federal do Rio Grande do Sul, 2007.
- [21] G. L. Kieling and R. M. Vicari, “Insertion of probabilistic knowledge into BDI agents construction modeled in bayesian networks,” in *Complex, Intelligent and Software Intensive Systems (CISIS), 2011 International Conference on*, pp. 115–122, IEEE, 2011.
- [22] D. G. Silva and J. C. Gluz, “AgentSpeak(PL): A new programming language for BDI agents with integrated bayesian network model,” in *Information Science and Applications (ICISA), 2011 International Conference on*, pp. 1–7, IEEE, 2011.
- [23] R. R. Bouckaert, E. Frank, M. A. Hall, G. Holmes, B. Pfahringer, P. Reutemann, and I. H. Witten, “WEKA — experiences with a java open-source project,” *The Journal of Machine Learning Research*, pp. 2533–2541, 2010.
- [24] I. H. Witten and E. Frank, *Data Mining: Practical machine learning tools and techniques*. Morgan Kaufmann, 2005.
- [25] R. Dimov, M. Feld, D. M. Kipp, D. A. Ndiaye, and D. D. Heckmann, “Weka: Practical machine learning tools and techniques with java implementations,” *AI Tools Seminar University of Saarland, WS*, vol. 6, no. 07, 2007.
- [26] M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I. H. Witten, “The WEKA data mining software: an update,” *ACM SIGKDD Explorations Newsletter*, vol. 11, no. 1, pp. 10–18, 2009.
- [27] S. Russell, *Artificial intelligence: A modern approach, 2/E*. Pearson Education India, 2003.
- [28] R. Socher, C. C. Lin, A. Ng, and C. Manning, “Parsing natural scenes and natural language with recursive neural networks,” in *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, pp. 129–136, 2011.
- [29] G. E. Hinton, “Learning multiple layers of representation,” *Trends in cognitive sciences*, vol. 11, no. 10, pp. 428–434, 2007.
- [30] R. Salakhutdinov and G. Hinton, “An efficient learning procedure for deep boltzmann machines,” *Neural Computation*, vol. 24, no. 8, pp. 1967–2006, 2012.
- [31] I. Jolliffe, *Principal component analysis*. Wiley Online Library, 2005.